

In situ upgrade of quantum simulators to universal computers

Benjamin Dive,¹ Alexander Pitchford,² Florian Mintert,¹ and Daniel Burgarth²

¹*Department of Physics, Imperial College, SW7 2AZ London, UK*

²*Institute of Mathematics, Physics and Computer Science,
Aberystwyth University, SY23 3FL Aberystwyth, UK*

(Dated: January 9, 2017)

The ability to perform a universal set of logic gates on a quantum simulator would come close to upgrade it into a universal quantum computer [1, 2]. Knowing how to do this is very hard as it requires a precise knowledge of the simulator. In most cases, it also needs to be itself simulated on a classical computer as part of an optimal control algorithm [3, 4]. This generally can not be done efficiently for the very reason that quantum computers provide an advantage over classical ones. Here we use a simulator to discover how to implement a universal set of gates on itself without knowing the details of its own workings. The method is scalable for a series of examples and is a practical way of upgrading quantum simulators to computers, as well as opening up new possible architectures.

Recent and ongoing work on building large quantum systems is leading to simulators that are able to model physical phenomena, allowing questions about the underlying science to be answered [2, 5, 6]. These machines are a register of quantum particles each storing a quantum bit (qubit) of information in two internal states. The presence of interactions between these leads to dynamics that, by varying control parameters in the system Hamiltonian, can replicate the quantum behaviour of systems of interest. This, however, is less general than a quantum computer which is able to perform a universal set of logic gates on the qubits [1].

Provided some control parameters can be varied in time, it is in principle possible to do an arbitrary gate on a quantum many-body system such as a quantum simulator [7, 8]. Finding the right time-dependency however relies almost exclusively on numerical methods, especially when physical constraints on the control fields are taken into account [3, 4]. These require a very precise knowledge of the parameters of a system, a daunting task for a machine with a huge number of degrees of freedom. Furthermore, they are intractable on a classical computer if the simulator we want to solve the problem for is large enough to do something new.

We circumvent these problems by showing how well-known existing numerical methods can be translated to run *in situ* on the quantum simulator itself, as illustrated in Fig. (1). This is a scalable, bootstrapping scheme for performing a universal quantum computation, needing classical computational and experimental resources that grow only polynomially with the number of qubits. An adaptive approach to finding controls is naturally used in laboratory work and has been studied as both fine-tuning tool in small systems [9, 10] and a way of correcting parameter drift [11]. Using simulators as oracles for reaching quantum states was explored in [12], and the potential for a quantum speed up in general optimisation problems in [13]. We show that this general approach can be transformed into a large scale, constructive method that changes the way we approach the control of many-body quantum systems. Independently, and concurrently to




FIG. 1. A classical computer finds a control pulse which enables a quantum simulator to perform logic gates. It does this in an iterative process by applying a control pulse to the simulator and then improving it based on the result of measurements.

the submission of this paper, a similar approach as ours was developed and tested experimentally for quantum state preparation [14].

The model we consider is a quantum simulator consisting of n qubits with some interactions between them such that they are connected on a timescale much shorter than the decoherence time, and with the ability to do fast single qubit operations. These requirements are significant, but much easier than demanding direct control over two qubit operations, and correspond to the state-of-the-art in systems involving trapped ions [16–18], cold atoms [19, 20] or NMR [21–23]. In these systems there already exist quantum simulators powerful enough to do simulations, and satisfy our requirements, but are not currently usable as computers as it is not known how to perform logic gates on them [1, 2]. In the model we consider, the connectedness of the qubits and the ability to do fast




FIG. 2. Outline of the process used in optimal control, the two red steps in the middle are done *in situ* in our scheme. The starting point is an initial set of control parameters, which we generated randomly for our examples. The evolution of the system with these parameters is then calculated. On a classical computer this requires solving the time-dependent Schrödinger equation numerically for a model of the system, while in our scheme this is simply implementing the controls on the simulator. Evaluating the gate fidelity in the classical case is straightforward, but in our scheme it requires some form of tomography to measure it. We derived good bound on this gate fidelity, Eq. (1), that can be found efficiently. If this fidelity is above a threshold, the process terminates successfully, otherwise the control parameters are updated based on the existing and previous runs, and the process repeats. There are many ways to update the controls classically [15] that can also be used in this method. In our examples we used a steepest ascent method which requires the gradient of the fidelity in the control parameters to also be measured; we show that this can also be done efficiently.

single qubit operations guarantees that the two requirements of the optimisation scheme are satisfied: there exists a universal gate set that can be reached at short times [24], and process tomography can be performed [25]. While other systems satisfy these requirements and the approach detailed here would work, we focus on this model for clarity. As single qubit operations are assumed, the gates that controls are needed for are entangling ones, canonically the controlled-not (C-NOT) gate, which are vastly harder to perform using conventional methods.

The steps for finding such a gate in the *in situ* scheme are outlined in Fig. (2). These are the same as in classical optimal control except that the control parameters to be optimised are implemented on the quantum machine itself, rather than simulated on a classical computer. This eliminates the need to characterise and simulate the system, thereby removing the key elements that prevent classical optimal control theory from being efficient. The computational cost for the *in situ* scheme is the number of runs of simulator needed to find a control pulse, given by $N_{\text{tot}} = N_{\text{iters}} N_{\text{fids}} N_{\text{reps}} N_{\text{meas}}$ where: N_{iters} is the number of iterations where the scheme goes around the loop of Fig. (2), N_{fids} the number of fidelities that must be measured per iteration, N_{reps} the number of repeats of each fidelity measurement required to get the desired accuracy, and N_{meas} is the number of measurements per fidelity. Most of these depend on the underlying algorithm used and, as is usual in control theory, we do not have analytic expressions for them, so we present numerical evidence of how they scale with n in




FIG. 3. The gate fidelity, and its local estimator Eq. 1, are plotted as a function of iteration step for one complete run of the *in situ* optimisation scheme. The system is a five-qubit Ising chain where the target is a C-NOT gate on the first two qubits and identity on the others. The algorithm minimised the infidelity of the local estimator. The exact infidelity is plotted at each step for comparison. It is lower in every case, and highly correlated with the estimated infidelity, such that minimising the former also minimises the latter almost monotonically and the landscape remains trap free. Furthermore the difference between the two decreases rapidly as the infidelity approaches 0. This demonstrates the validity of maximising the local estimator of the fidelity as a proxy for maximising the true gate fidelity.

stead. The exception to this is the fidelity, where we have a good bound that can be measured efficiently. For the range of examples we considered, we found that the total cost, N_{tot} , scales polynomially as $O(n^4)$.

Measuring the gate fidelity between the (potentially noisy) realised dynamics and the target unitary is an important part of the protocol. The standard method of doing this is to perform a variant of process tomography, known as certification, which requires a number of measurements that scales exponentially: $N_{\text{meas}} = O(d^2) = O(2^{2n})$ [27]. However for cases of interest where the target gate has a tensor product structure, $U = \bigotimes U_i = \text{C-NOT}_1 \otimes \mathbb{1}_2 \otimes \mathbb{1}_3 \otimes \dots$, it is possible to get around this. The fidelity over the whole system is bounded by the fidelities of the tensor subsystems according to

$$F(M, U) \geq F_{LE}(M, U) = 1 - \sum_i (1 - F(M_i, U_i)) \quad (1)$$

where $F(M, U)$ is the gate fidelity [28] between the evolution M and target unitary U , the sum is over all elements of the tensor product, and $M_i(\rho_i) = M(\rho_i \bigotimes_{j \neq i} \frac{1}{d_j} \mathbb{1}_j)$ is the reduced dynamical map acting on any state in subsystem i and the maximally mixed state everywhere else, which is proved in the appendix following methods from [29]. By using this local estimator of the fidelity F_{LE} as the figure of merit to optimise over, the true fidelity is guaranteed to be at least as high and, as we can see in Fig. (3), converges well. Importantly, it can be measured

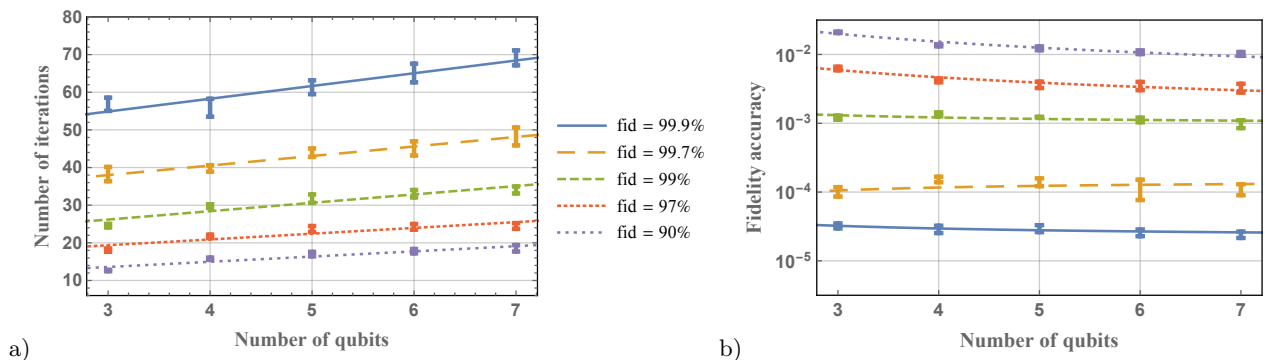


FIG. 4. **Numerical simulation of the experimental cost of finding a C-NOT gate in an Ising chain using steepest ascent *in situ*.** The number of iterations and the accuracy required for the optimisation protocol to succeed is plotted for a chain of qubits with nearest neighbour interactions $\sigma_z \otimes \sigma_z$ evolving for a time π (for units of $\hbar = 1$) with σ_x and σ_y controls at each site. In figure a) the fidelity accuracy is picked to give a 50% success rate and the number of iterations plotted for chains of various length and for different target fidelities (as these are F_{LE} , the true gate fidelity will be somewhat better). We see a strong linear relation in the number of iterations required as a function of the number of qubits giving $N_{\text{iters}} = O(n)$. The required precision to achieve 50% success rate irrespective of the number of iterations is shown in b) for the same system, and matches very well with a theoretical fit of $O(1/n)$. This fit arises by noting that to reach a gate infidelity of ϵ ought to require measuring the fidelity to an accuracy $O(\epsilon)$; as this is calculated from the sum of the fidelities of the subsystem, it is reasonable that these need to be measured to an accuracy $O(\epsilon/n)$. This argument is backed up by how well the data lies on the curves. From this data, we can estimate N_{reps} and N_{fids} . The required accuracy scales as $O(1/n)$ which, with the central limit theorems, implies that the number of measurement repetitions N_{reps} scales as $O(n^2)$. The number of fidelity measurements per iteration of the control parameters, N_{fids} , scales as $O(n)$. This is because a steepest ascent algorithm requires the gradient of the fidelity for each of the control parameters in order to update them. The most direct way of obtaining these gradients is to repeat the simulation with each of the control parameters increased by a small finite amount in turn and calculate the approximate finite difference gradient (this was not done in our numerical simulations as it was too large a computational cost, an analytic gradient obtained using GRAPE [26] was used instead). This requires one additional fidelity to be measured per control parameter. In our algorithm (and GRAPE), each of the control Hamiltonians is divided into a number of time slots of constant amplitude. Surprisingly, we found that the required number of time slots for the protocol to work on the Ising chain was 12, independent of the length of the chain. As there are 2 control parameters per qubit (corresponding to σ_x and σ_y) per time slot, this gives $N_{\text{fids}} = O(n)$.

efficiently due to the need to perform certification over small subspaces only, as scaling the scheme up increases the number of subspaces but not their size. This leads to, $N_{\text{meas}} = O(1)$, an exponential decrease in the cost of estimating the gate fidelity.

The scaling relation of the other components of N_{tot} depends on the underlying classical algorithm. We used a steepest ascent gradient method similar to the classical Gradient Ascent Pulse Engineering (GRAPE) algorithm [26] which is commonly used in quantum control with great success. With extensive numerical optimisation and simulations on a Python platform that we expanded [30, 31] we find in Fig. (4) that to reach a C-NOT gate in an Ising chain of n qubits the number of runs of the simulation required is $N_{\text{tot}} = O(n^4)$. This efficient scaling is verified for systems of up to 7 qubits numerically. Extending this to much larger systems is implausible due to the difficulty of simulating and optimising quantum systems on a classical computers; the current numerical data took the equivalent of 2 years of laptop computational time. Indeed, finding a way around this is the very advantage of the *in situ* optimisation scheme; testing this scaling relation for large systems would be an interesting experiment that could only be done on a quantum

machine. We also surveyed a range of different possible simulator topologies and interaction type, as summarised in Fig. (5). In all cases the scheme worked at reasonable cost, with the Ising chain and star working particularly well. A last point to investigate is that the scheme might fail and not converge to the desired gate. As we are using a gradient based method, this would happen if there is a local maximum in the fidelity landscape where the algorithm gets stuck. Numerically we did not find this and the protocol converged all the time provided the fidelity was calculated to a high enough accuracy (this required accuracy is investigated in Fig. (4)). The lack of traps is supported by recent results which state that, for well behaved fidelity functions, a generic control system will have no such traps [32, 33].

A future direction to take this work is to apply it to another important aspect of quantum computation, error correction. The protocol detailed here can be used in much the same way for this by replacing the target operation from a C-NOT gate, to one protecting some logical qubits. Preliminary results show that with a tuneable interaction the system can discover decoherence free subspaces and simple error correcting codes this way. Work remains on what the right tasks to ask of the simulator

Topology	Coupling	T	N_{iters}	N_{fids}
chain	Ising	π	60	120
star	Ising	π	214	120
fully connected	Ising	12π	295	1600
chain	Heisenberg	16π	585	1600
star	Heisenberg	12π	1043	1600
fully connected	Heisenberg	12π	881	1600

FIG. 5. The cost of performing the *in situ* optimisation scheme is investigated for a range of different 5 qubit systems using the same numerical method as in Fig. (4). The differences between the systems are their topology (a linear chain with nearest neighbour interactions, a star where all interact with a central qubit only, or fully connected where the interaction strengths are also randomised) and the interaction type (Ising with $\sigma_z \otimes \sigma_z$, or Heisenberg with $\sigma_x \otimes \sigma_x + \sigma_y \otimes \sigma_y + \sigma_z \otimes \sigma_z$). In each case the controls are σ_x and σ_y on the single qubits, and the target operation is a C-NOT gate on two qubits and identity on the rest. N_{iters} is the average number of iterations required to achieve the 99.9% fidelity target. T is the time allowed for the system to evolve, N_{reps} and N_{fids} are as in the main text and Fig. (4). We see that for all six systems the protocol can find the desired entangling gate, and does so at reasonable experimental cost. This indicates that the approach works for a wide range of possible quantum simulators. The case of the Ising chain and star are particularly easy for the scheme requiring a low evolution time, number of iterations, and number of fidelity measurements per iteration.

are, and on showing the scalability of this approach.

The *in situ* optimal control scheme introduced here runs a classical optimisation algorithm on a quantum simulator in order to turn it into a universal quantum computer. While the underlying algorithm we used was gradient based, many different ones could be used instead to impose different constraints on the shape of the control pulse. All the evidence points towards the cost of doing this scaling polynomially with the size of the simulator, making such an approach feasible for large quantum systems. A key aspect of this is to bound the gate fidelity of the whole system by the gate fidelities of single and two qubit blocks, which can be measured efficiently. As well as upgrading quantum simulators into computer, it suggests that it is not necessary to restrict architectures of quantum computers to those where controls can be found analytically: having some interactions and local controls is enough to efficiently bootstrap a quantum machine into a universal computer.

This work was supported by EPSRC through the Quantum Controlled Dynamics Centre for Doctoral Training, the EPSRC Grant No. EP/M01634X/1, and the ERC Project ODYCQUENT. We are grateful to HPC Wales for giving access to the cluster that was used to perform the numerical simulations. Many thanks to Stephen Glaser and David Leiner for discussions on possible implementations.

-
- [1] D. DiVincenzo, arXiv:quant-ph/0002077 (2000).
[2] T. H. Johnson, S. R. Clark, and D. Jaksch, EPJ Quantum Technol. **1**, 1 (2014).
[3] S. Machnes, U. Sander, S. J. Glaser, P. de Fouquières, A. Gruslys, S. Schirmer, and T. Schulte-Herbrüggen, Phys. Rev. A **84**, 022305 (2011).
[4] S. J. Glaser, U. Boschain, T. Calarco, C. P. Koch, W. Köckenberger, R. Kosloff, I. Kuprov, B. Luy, S. Schirmer, T. Schulte-Herbrüggen, D. Sugny, and F. K. Wilhelm, Eur. Phys. J. D **69** (2015).
[5] J. I. Cirac and P. Zoller, Nat. Phys. **8**, 264 (2012).
[6] I. M. Georgescu, S. Ashhab, and F. Nori, Rev. Mod. Phys. **86**, 153 (2014).
[7] S. Lloyd, Phys. Rev. Lett. **75**, 346 (1995).
[8] F. Albertini and D. D'Alessandro, Linear Algebra Appl. **350**, 213 (2002).
[9] D. J. Egger and F. K. Wilhelm, Phys. Rev. Lett. **112**, 240503 (2014).
[10] J. Kelly, R. Barends, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, a. G. Fowler, I. C. Hoi, E. Jeffrey, A. Megrant, J. Mutus, C. Neill, P. J. J. O'Malley, C. Quintana, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, a. N. Cleland, and J. M. Martinis, Phys. Rev. Lett. **112**, 240504 (2014).
[11] J. Kelly, R. Barends, A. G. Fowler, A. Megrant, E. Jeffrey, T. C. White, D. Sank, J. Y. Mutus, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, E. Lucero, M. Neeley, C. Neill, P. J. J. O'Malley, C. Quintana, P. Roushan, A. Vainsencher, J. Wenner, and J. M. Martinis, Phys. Rev. A **94**, 032321 (2016).
[12] J. Li, X. Yang, X. Peng, and C.-P. Sun, arXiv:1608.00677 (2016).
[13] P. Rebentrost, M. Schuld, F. Petruccione, and S. Lloyd, arxiv:1612.01789 (2016).
[14] D. Lu, K. Li, J. Li, H. Katiyar, A. J. Park, G. Feng, T. Xin, H. Li, G. Long, A. Brodutch, J. Baugh, B. Zeng, and R. Laflamme, arxiv:1701.01198 (2017).
[15] D. Dong and I. R. Petersen, IET Control Theory Appl. **4**, 2651 (2010).
[16] M. Johanning, A. F. Varón, and C. Wunderlich, J. Phys. B At. Mol. Opt. Phys. **42**, 154009 (2009).
[17] B. P. Lanyon, C. Hempel, D. Nigg, M. Müller, R. Geritsma, F. Zähringer, P. Schindler, J. T. Barreiro, M. Rambach, G. Kirchmair, M. Hennrich, and P. Zoller, Science (80-.). **334**, 57 (2011).
[18] R. Blatt and C. F. Roos, Nat. Phys. **8**, 277 (2012).
[19] I. Bloch, J. Dalibard, and S. Nascimbène, Nat. Phys. **8**, 267 (2012).
[20] H. Labuhn, D. Barredo, S. Ravets, S. de Léséleuc, T. Macrì, T. Lahaye, and A. Browaeys, Nature **534**, 667 (2016).
[21] X. H. Peng and D. Suter, Front. Phys. China **5**, 1 (2010).
[22] J. Cai, A. Retzker, F. Jelezko, and M. B. Plenio, Nat. Phys. **9**, 168 (2013).
[23] J. Geng, Y. Wu, X. Wang, K. Xu, F. Shi, Y. Xie, X. Rong, and J. Du, Phys. Rev. Lett. **117**, 170501 (2016).

- [24] J. L. Dodd, M. A. Nielsen, M. J. Bremner, and R. T. Thew, Phys. Rev. A **65**, 040301 (2002).
- [25] J. F. Poyatos, J. I. Cirac, and P. Zoller, Phys. Rev. Lett. **78**, 390 (1997).
- [26] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, J. Magn. Reson. **172**, 296 (2005).
- [27] M. P. Da Silva, O. Landon-Cardinal, and D. Poulin, Phys. Rev. Lett. **107**, 210404 (2011).
- [28] A. Gilchrist, N. K. Langford, and M. a. Nielsen, Phys. Rev. A **71**, 062310 (2005).
- [29] M. Cramer, M. Plenio, S. Flammia, R. Somma, D. Gross, S. Bartlett, O. Landon-Cardinal, D. Poulin, and Y. Liu, Nat. Commun. **1**, 149 (2010).
- [30] J. R. Johansson, P. D. Nation, and F. Nori, Comput. Phys. Commun. **184**, 1234 (2013).
- [31] R. Johansson, P. Nation, A. Pitchford, C. Granade, A. L. Grimsmo, Markusbaden, A. Vardhan, P. Migdal, Kafischer, D. Vasilyev, B. Criger, J. Feist, Alexbr, F. Henneke, D. Meiser, F. Ziem, Nwlambert, S. Krashtanov, R. Brierley, R. Heeres, J. Hörsch, Qi, S. Whalen, N. Tezak, J. Neergaard-Nielsen, E. Hontz, Amit, and Adriaan, “QuTiP-4.0.0,” (2016).
- [32] T. S. Ho, J. Dominy, and H. Rabitz, Phys. Rev. A - At. Mol. Opt. Phys. **79**, 013422 (2009).
- [33] B. Russell, H. Rabitz, and R. Wu, arxiv:1608.06198 (2016).

APPENDIX - LOCAL ESTIMATOR OF THE FIDELITY

We derive a bound for the gate fidelity $F(M, U) = \langle \psi | \rho_M | \psi \rangle$ where $|\psi\rangle\langle\psi| = U \otimes id |\Omega\rangle\langle\Omega|$ and $\rho = M \otimes id |\Omega\rangle\langle\Omega|$ are the Choi states of U and M respectively, given by those expressions, where id is the identity map and $|\Omega\rangle = \sum_k |kk\rangle$ is a maximally entangled state between the original and doubled space. We consider the case where the target operation U is unitary and has a tensor product structure such that $U = \bigotimes_i U_i$. The Choi state of U has the same structure $|\psi\rangle = \bigotimes_i |\psi_i\rangle$ where each $|\psi_i\rangle$ is on the doubled Hilbert space of U_i .

To find a bound for $F(M, U)$ we begin by constructing the projectors $h_i = \mathbb{1}_i - |\psi_i\rangle\langle\psi_i|$ for each tensor component of U . These have a single 0 eigenvalue, and the others are 1. These projectors are grouped together to

from a Hamiltonian $H = \sum_i h_i \otimes \mathbb{1}_{\bar{i}}$ such that each h_i acts on its own part of the Hilbert space and is identity on the rest. This has a single 0 eigenvalue E_0 with eigenstate $|\psi\rangle$, while all its other eigenvalues are positive integers. By expanding this Hamiltonian in its eigenbasis we have

$$\begin{aligned} \text{Tr}[H\rho] &= \sum_{k \geq 0} E_k \langle E_k | \rho | E_k \rangle \\ &\geq \sum_{k > 0} \langle E_k | \rho | E_k \rangle \\ &\geq 1 - \langle \psi | \rho | \psi \rangle \\ F(M, U) &= \langle \psi | \rho | \psi \rangle \geq 1 - \text{Tr}[H\rho] \end{aligned}$$

The expectation value of the Hamiltonian can be evaluated according to

$$\begin{aligned} \text{Tr}[H\rho] &= \sum_i \text{Tr}[(\mathbb{1}_i - |\psi_i\rangle\langle\psi_i|)\rho] \\ &= \sum_i (1 - \langle \psi_i | \rho_i | \psi_i \rangle), \end{aligned}$$

where $\rho_i = \text{Tr}_i[\rho]$. This is also the Choi state of the map $M_i(\cdot) = M(\cdot \otimes_{j \neq i} \frac{1}{d_j} \mathbb{1}_j)$, which results in Eq. (1)

$$F(M, U) \geq 1 - \sum_i (1 - F(M_i, U_i)).$$

One way of measuring F_{LE} is to prepare the state of one subspace in a basis state and all the others in a maximally mixed state, perform the simulation, measure the initial subspace, and then repeat for a tomographically complete basis set and for each subsystem; giving a cost of $N_{\text{meas}} = \sum_i O((d_i)^2) = O(n)$. However by noting that a maximally mixed state is a random mixture of pure states, the fidelity of each subsystem can be measured at the same time by preparing each one in a random pure basis state. In this case there is no scaling with the number of qubits as the number of repetitions required depends only on the size of the largest subsystem, which gives $N_{\text{meas}} = O(1)$.